

Heterogeneous Gain Learning and Long Swings in Asset Prices

Blake LeBaron *

International Business School

Brandeis University

November 15, 2010

Abstract

This paper considers the impact of heterogeneous gain learning in an asset pricing model. A relatively stylized model is shown to generate persistent swings of asset prices from their fundamental values which replicates long range samples of U.S financial data. The detailed mechanisms of the learning models are then explored. Evidence suggests that agents' perceptions of risk and its dynamics and persistence are important in generating appropriate price/fundamental dynamics. Agents putting large amounts of weight on the recent past in their volatility models control a large fraction of wealth, and are important in perpetuating the volatility magnifying dynamics of the market.

Keywords: Learning, Asset Pricing, Financial Time Series, Evolution, Memory

*International Business School, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a research associate at the National Bureau of Economic Research.

1 Introduction

Many asset prices exhibit large and highly persistent deviations from their fundamental values yielding potential long run predictability as well as one of the major puzzles of financial economics.¹ The fact that asset prices can move from simple benchmark rational pricing levels, and then stay far from these levels for some time is a major puzzle. The swings can be both short or long in duration, and their time series show surprisingly few regular patterns when one analyzes their long range behavior. Many explanations have been proposed involving varying levels of rationality, knowledge, and learning, but no one explanation has dominated the debate. This paper explores an under parameterized learning model with heterogeneous gain parameters and traders using differing perspectives on history.²

Gain parameters are present in almost all learning models, and they control the weight given to new observations as learners update their parameters. They can be thought of as controlling how much data from the past agents use to estimate new forecast parameters. They also can be interpreted as related to the signal to noise ratio in an unobserved state variable world.³ Learning models can also entertain both decreasing and constant gain levels. In decreasing gain models, the weight on new observations approaches zero as time passes, often generating useful convergence results. Constant gain learners always are forgetting some data from the distant past, and they continue to give new data points the same weight as they come in. Committing to a fixed gain level locks agents onto a certain belief about the overall stationarity of the time series they are observing. Entertaining various beliefs about gain parameters would be suggested in a world in which priors about data stationarity, persistence, and long range dynamic features may be spread over many possibilities. Also, the data series may be short enough, and the features long enough, that the data alone does not lead to much convergence in these beliefs. In other words, the belief that we are in a “new world,” and the last 5 years of data are all that is relevant will always hold some appeal in such a world, and statistical evidence firmly rejecting this idea may be very weak.

It is not hard to find candidate time series that appear very persistent, but the persistent behavior is often irregular and appears at horizons that are long relative to the sample length. This paper will concentrate specifically on price/dividend ratios from U.S. stock markets,

¹ The early evidence is in Shiller (1981) and Campbell & Shiller (1988). For a textbook treatment see Campbell, Lo & MacKinlay (1996). A recent survey on this literature is contained in Lettau & Ludvigson (2010).

² The area of learning models and asset pricing is very large. A recent survey by Pastor & Veronesi (2009) is a good starting place. A more general survey covering the broader macro economic context would be Evans & Honkapohja (2008). All of these are influenced by the early work collected in Phelps (1970) and Frydman & Phelps (1983). Also, Timmerman (1993), and Hong & Stein (1999) are early examples of learning models capable of generating market volatility which will be a key aspect of the model presented here. Frydman & Goldberg (2007) is another summary of the adaptive learning approach, that also presents models which display long persistent deviations in foreign exchange rates from PPP benchmarks. A recent model with Bayesian learning agents which replicates movements in long run price/dividend ratios is Adam & Marcet (2010a).

³In the traditional Kalman filter the gain parameter is explicitly defined by the structure of the model.

but there are many others. Inflation rates, interest rates, and real exchange rates, are all possible candidates for this. More subtly, the levels of volatility in financial markets also demonstrate this form of “extreme persistence.” In all cases an argument could be made for some very persistent process that could formally be a long memory, or fractionally integrated system, or their may simply be infrequent changes in regime driving results. In a learning environment agents may be confused between all these possibilities both because definitive statistical tests don’t exist, or sample sizes may simply be too short.

This paper borrows from many parts of the learning literature. At the core agents will be using systems derived from adaptive expectations and recursive least squares algorithms.⁴ Agents chose from a set of algorithms all estimated using varying gain levels. Agents also choose from a discrete set of possible forecasting rules designed to span the space of reasonable strategies used in markets.⁵ Finally, the behavior of the market depends on the distribution of wealth across strategies. This is driven both by the active selection of forecast rules by agents, as well as by the passive movement of wealth toward successful strategies.⁶

The importance of heterogeneous gain levels in a financial market was shown in the early computational models such as LeBaron, Arthur & Palmer (1999). In that paper, convergence of the market to a rational expectations equilibrium depended critically on how fast agents were updating their forecasting models. Frequent updates, which forced agents to use relatively short time series for decision making, led to a rich pattern in prices and volume, far from any recognizable equilibrium. However, restricting updates to occur relatively infrequently, and forcing agents to use long time series to evaluate rules, caused the market to converge to a well defined homogeneous rational expectations equilibrium.⁷

The macro economic environment used here will be simple, and easily recognizable. It is a form of “Lucas tree” model of asset prices as in Lucas (1978). A single asset pays a stochastic dividend calibrated to the trend and volatility of U.S. real dividend series. This model is too simple to fit all the facts and puzzles from the macro finance world. However,

⁴ A complete treatment on this topic is contained in Evans & Honkapohja (2001). An exploration of many forms of heterogeneity in adaptive learning models, including gain parameters, is in Honkapohja & Mitra (2006).

⁵The model shares much basic intuition in this area with the early work in Brock & Hommes (1998), Day & Huang (1990), Frankel & Froot (1988), Kirman (1991), and Lux (1998). Recent surveys to this literature can be found in Chiarella, Dieci & He (2009) and Hommes & Wagener (2009), and Tesfatsion & Judd (2006).

⁶This overlaps with the large literature on growth optimal portfolio strategies which can be traced back to Kelley (1956), and a more recent treatment can be found in Blume & Easley (1990). A recent survey on this is Evstigneev, Hens & Schenk-Hoppe (2009).

⁷ Similar results, but in very different models can be found in LeBaron (2006), and in Thurner, Dockner & Gaunersdorfer (2002). This paper directly considers the destabilizing impact of large gain, or “short memory” traders on market dynamics. Similar questions about how much data agents should be using from the past are considered in Mitra (2005). A recent model stressing what might happen when agents overweight shorter run trends, ignoring longer length reversals, is Fuster, Laiso & Mendel (2010). Hommes (2010) surveys growing experimental literature which often shows people following short term trends when reporting their expectations in controlled laboratory settings. Finally, Dacorogna, Gencay, Muller, Olsen & Pictet (2001) presents a philosophy, and some time series models, for markets populated by agents with many different time perspectives.

it seems to be in a good middle ground, allowing for some quantitative comparisons of the learning model to actual time series, but not getting so complex as to hide basic features which are central to the learning dynamics. For example, macro level shocks should always be viewed as extremely costly in terms of using up modeling degrees of freedom. This model is restricted to only one aggregate shock coming from the movement in aggregate dividends. This restriction is critical, and makes it easier to interpret how much of the price volatility from the model can be interpreted as an endogenous feature of learning, as opposed to coming from other poorly specific macro level shocks.⁸

A final interesting aspect to learning in this model is the explicit consideration of risk and return. This was used in LeBaron et al. (1999). More recently several authors have begun exploring this issue as in Branch & Evans (2008). This allows separation of the dynamics into return and volatility forecasts. The functional style of return forecasts is allowed to vary across forecast families, but the volatility forecasts all follow a common form. Large returns are interpreted as increases in risk by most agents. This distinction seems plausible, and is critical in how the market is able to generate price fundamental deviations.

Section 2 gives a short explanation of the basic model. Section 3.1 examines some benchmark simulation runs, and observes the output of the model as compared to actual financial times series. This section cannot report all the results from the model, since this would distract the paper from its main mission of exploring mostly long range swings in prices from their fundamentals. Sections 3.2 and 3.3 look at many internal mechanism of the agents and forecasts in use, and how wealth moves across these over time. Learning algorithms appear to be behaving in a predictable fashion, and interesting dynamics come from how agent wealth selects rules over time. The final section concludes and introduces some questions for researchers working on learning in financial markets.

2 Model description

This section will briefly describe the structure of the model. It combines components from many parts of the learning literature. The goal is to build a heterogeneous agent asset market which is as parsimonious as possible, but can still do a good job of replicating financial market features. Also, its inner workings should be simple enough for detailed analysis, meet a general plausibility test, and be rich enough to understand several aspects of how wealth moves around in a learning investment environment.

⁸The model requires a variety of independent micro shocks, but these should be viewed as much more plausible, and less costly, than adding further aggregate shocks.

2.1 Assets

The market consists of only two assets. First, there is a risky asset paying a stochastic dividend,

$$d_{t+1} = d_g + d_t + \epsilon_t, \quad (1)$$

where d_t is the log of the dividend paid at time t . Time will be incremented in units of weeks. Lower case variables will represent logs of the corresponding variables, so the actual dividend is given by,

$$D_t = e^{d_t}. \quad (2)$$

The shocks to dividends are given by ϵ_t which is independent over time, and follows a Gaussian distribution with zero mean, and variance, σ_d^2 , that will be calibrated to actual long run dividends from the U.S. The dividend growth rate would then be given by $e^{g+(1/2)\sigma_d^2}$ which is approximately $D_g = d_g + (1/2)\sigma_d^2$.

The return on the stock with dividend at date t is given by

$$R_t = \frac{P_t + D_t - P_{t-1}}{P_{t-1}}, \quad (3)$$

where P_t is the price of the stock at time t . Timing in the market is critical. Dividends are paid at the beginning of time period t . Both P_t and D_t are part of the information set used in forecasting future returns, R_{t+1} . There are I individual agents in the model indexed by i . The total supply of shares is fixed, and set to unity,

$$\sum_{i=1}^I S_{t,i} = 1. \quad (4)$$

There is also a risk free asset that is available in infinite supply, with agent i holding $B_{t,i}$ units at time t . The risk free asset pays a rate of R_f which will be assumed to be zero in most simulations. This is done for two important reasons. It limits the injection of outside resources to the dividend process only. Also, it allows for an interpretation of this as a model with a perfectly storable consumption good along with the risky asset. The standard intertemporal budget constraint holds for each agent i ,

$$W_{t,i} = P_t S_{t,i} + B_{t,i} + C_{t,i} = (P_t + D_t) S_{t-1,i} + (1 + R_f) B_{t-1,i}, \quad (5)$$

where $W_{t,i}$ represents the wealth at time t for agent i .

2.2 Preferences

Portfolio choices in the model are determined by a simple myopic power utility function in future wealth. The agent's portfolio problem corresponds to,

$$\max_{\alpha_{t,i}} \frac{E_t^i W_{t+1,i}^{1-\gamma}}{1-\gamma}, \quad (6)$$

$$st. \quad W_{t+1,i} = (1 + R_{t+1,i}^p)(W_{t,i} - C_{t,i}), \quad (7)$$

$$R_{t+1,i}^p = \alpha_{t,i} R_{t+1} + (1 - \alpha_{t,i}) R_f. \quad (8)$$

$\alpha_{t,i}$ represents agent i 's fraction of savings ($W - C$) in the risky asset. It is well known that the solution to this problem yields an optimal portfolio weight given by,

$$\alpha_{t,i} = \frac{E_t^i(r_{t+1}) - r_f + \frac{1}{2}\sigma_{t,i}^2}{\gamma\sigma_{t,i}^2} + \epsilon_{t,i}, \quad (9)$$

with $r_t = \log(1 + R_t)$, $r_f = \log(1 + R_f)$, $\sigma_{t,i}^2$ is agent i 's estimate of the conditional variance at time t , and $\epsilon_{t,i}$ is an individual shock designed to make sure that there is some small amount of heterogeneity to keep trade operating.⁹ It is distributed normally with variance, σ_ϵ^2 .

In the current version of the model neither leverage nor short sales are allowed. The fractional demand is restricted to $\alpha_{t,i}$ to $\alpha_L \leq \alpha_{t,i} \leq \alpha_H$. The addition of both these features is important, but adds significant model complexity. One key problem is that with either one of these, one must address problems of agent bankruptcy, and borrowing constraints. Both of these are not trivial, and involve many possible implementation details.

Consumption will be assumed to be a constant fraction of wealth, λ . This is identical over agents, and constant over time. The intertemporal budget constraint is therefore given by

$$W_{t+1,i} = (1 + R_{t+1,i}^p)(1 - \lambda)W_{t,i}. \quad (10)$$

This also gives the current period budget constraint,

$$P_t S_{t,i} + B_{t,i} = (1 - \lambda)((P_t + D_t)S_{t-1,i} + (1 + R_f)B_{t-1,i}). \quad (11)$$

This simplified portfolio strategy will be used throughout the paper. It is important to note that the fixed consumption/wealth, myopic strategy approach given here would be optimal in a standard intertemporal model for consumption portfolio choice subject to two key assumptions. First, the intertemporal elasticity of substitution would have to be unity

⁹The derivation of this follows Campbell & Viceira (2002). It involves taking a Taylor series approximation for the log portfolio return.

to fix the consumption wealth ratio, and second, the correlation between unexpected returns and certain state variables would have to be zero to eliminate the demand for intertemporal hedging.¹⁰

2.3 Expected Return Forecasts

The basic problem faced by agents is to forecast both expected returns and the conditional variance one period into the future. This section will describe the forecasting tools used for expected returns. A forecast strategy, indexed by j , is a method for generating an expected return forecast $E^j(r_{t+1})$. Agents, indexed by i , will choose forecast rules, indexed by j , according to an expected utility objectives.

All the forecasts will use long range forecasts of expected values using constant gain learning algorithms equipped with the minimum gain level, denoted g_L .

$$\bar{r}_t = (1 - g_L)\bar{r}_{t-1} + g_L r_t \quad (12)$$

$$\bar{pd}_t = (1 - g_L)\bar{pd}_{t-1} + g_L pd_{t-1} \quad (13)$$

$$\bar{\sigma}_t^2 = (1 - g_L)\bar{\sigma}_{t-1}^2 + g_L(r_t - \bar{r}_t)^2 \quad (14)$$

$$\bar{\sigma}_{pd,t}^2 = (1 - g_L)\bar{\sigma}_{pd,t-1}^2 + g_L(pd_t - \bar{pd}_t)^2 \quad (15)$$

The forecasts used will combine four linear forecasts drawn from well known forecast families. The first of these is a form of adaptive expectations which corresponds to,

$$f_t^j = f_{t-1}^j + g_j(r_t - f_{t-1}^j). \quad (16)$$

Forecasts of expected returns are dynamically adjusted based on the latest forecast error. This forecast format is simple and generic. It has roots connected to adaptive expectations, trend following technical trading, and also Kalman filtering.¹¹ The critical parameter is the gain level represented by g_j . This determines the weight that agents put on recent returns and how this impacts their expectations of the future. Forecasts with a large range of gain parameters will compete against each other in the market. Finally, this forecast will be trimmed in that it is restricted to stay between the values of $[-h_j, h_j]$. These will be set to relatively large values, and are randomly distributed across rules.

¹⁰See Campbell & Viceira (1999) for the basic framework. Also, see Giovannini & Weil (1989) for early work on determining conditions for myopic portfolio decisions. Hedging demands would only impose a constant shift on the optimal portfolio, so it is an interesting question how much of an impact this might have on the results.

¹¹See Cagan (1956), Friedman (1956), Muth (1960), and Phelps (1967) for original applications. A nice summary of the connections between Kalman filtering, adaptive expectations, and recursive least squares is given in Sargent (1999). A recent paper demonstrating a more general connection to state space models and expected returns is Pastor & Stambaugh (2009).

The second forecasting rule is based on a classic fundamental strategy. This forecast uses log price dividend ratio regressions as a basis for forecasting future returns,

$$f_t^j = \bar{r}_t + \beta_t^j(pd_t - \bar{pd}_t). \quad (17)$$

where pd_t is $\log(P_t/D_t)$. Although agents are only interested in the one period ahead forecasts the P/D regressions will be estimated using the mean return over the next M_{PD} periods, where $M_{PD} = 52$ weeks.

The third forecast rule will be based on simple linear regressions. It is a predictor of returns at time t given by

$$f_t^j = \bar{r}_t + \sum_{i=0}^{M_{AR}-1} \beta_{t,i}^j(r_{t-i} - \bar{r}_t) \quad (18)$$

This strategy works to eliminate short range autocorrelations in returns series through its behavior, and $M_{AR} = 3$ for all runs in this paper. It will be referred to as the Short AR forecast.

The previous two rules will be estimated each period using recursive least squares. There are many examples of this for financial market learning.¹² The key difference is that this model will stress heterogeneity in the learning algorithms with wealth shifting across many different rules, each using a different gain parameter in its online updating.¹³

The final rule is a buy and hold strategy using the long run mean, \bar{r}_t , for the expected return, and the long run variance, $\bar{\sigma}_t^2$, as the variance estimate. This portfolio fraction is then determined by the demand equation used by the other forecasting rules. This gives a useful passive benchmark strategy which can be monitored for relative wealth accumulation in comparison with the other active strategies.

2.4 Regression Updates

Forecasting rules are continually updated. The adaptive forecast only involves fixed forecast parameters, so its updates are trivial, requiring only the recent return. The two regression forecasts are updated each period using recursive least squares.

All the rules assume a constant gain parameter, but each rule in the family corresponds to a different gain level. This again corresponds to varying weights for the forecasts looking

¹² See Evans & Honkapohja (2001) for many examples, and also very extensive descriptions of recursive least squares learning methods.

¹³ Another recent model stressing heterogeneity in an OLS learning environment is Georges (2008) in which OLS learning rules are updated asynchronously.

at past data. The fundamental regression is run using the long range return,

$$\tilde{r}_t = \frac{1}{M_{PD}} \sum_{j=1}^{M_{PD}} r_{t-j+1} \quad (19)$$

The fundamental regression is updated according to,

$$\begin{aligned} \beta_{t+1}^j &= \beta_t^j + \frac{g_j}{\sigma_{pd,t}^2} p d_{t-M_{PD}} u_{t,j} \\ u_{t,j} &= (\tilde{r}_t - f_{j,t-M_{PD}}) \end{aligned} \quad (20)$$

Also, β_t^j is restricted to be between 0 and -0.05 . The zero upper bound on β makes sure this strategy is mean reverting, with an overall stabilizing impact on the market.

For the lagged return regression this would be,

$$\begin{aligned} \beta_{t+1,i}^j &= \beta_{t,i}^j + \frac{g_j}{\sigma_{r,t}^2} r_{t-i} u_{t,j}, \\ u_{t,j} &= (r_t - f_t^j) \end{aligned} \quad (21)$$

where g_j is again the critical gain parameter, and it varies across forecast rules.¹⁴ In both forecast regressions the forecast error, $u_{t,j}$, is trimmed. If $u_{t,j} > h_j$ it is set to h_j , and if $u_{t,j} < -h_j$ it is set to $-h_j$. This dampens the impact of large price moves on the forecast estimation process.

2.5 Variance Forecasts

The optimal portfolio choice demands a forecast of the conditional variance as well as the conditional mean.¹⁵ The variance forecasts will be generated from adaptive expectations as in,

$$\hat{\sigma}_{t+1,j}^2 = \hat{\sigma}_{t,j}^2 + g_{j,\sigma} (e_{t,j}^2 - \hat{\sigma}_{t,j}^2) \quad (22)$$

$$e_{t,j}^2 = (r_t - f_{t-1}^j)^2, \quad (23)$$

where $e_{t,j}^2$ is the squared forecast error at time t , for rule j . The above conditional variance estimate is used for all the rules. There is no attempt to develop a wide range of variance forecasting rules, reflecting the fact that while there may be many ways to estimate a

¹⁴This format for multivariate updating is only an approximation to the true recursive estimation procedure. It is assuming that the variance/covariance matrix of returns is diagonal. Generated returns in the model are close to uncorrelated, so this approximation is probably reasonable. This is done to avoid performing many costly matrix inversions.

¹⁵Several other papers have explored the dynamics of risk and return forecasting. This includes Branch & Evans (2008) and Gaunersdorfer (2000). In LeBaron (2001a) risk is implicitly considered through the utility function and portfolio returns. Obviously, methods that parameterize risk in the variance may miss other components of the return distribution that agents care about, but the gain in tractability is important.

conditional variance, they often produce similar results.¹⁶ This forecast method has many useful characteristics as a benchmark forecast. First, it is essentially an adaptive expectations forecast on second moments, and therefore shares a functional form similar to that for the adaptive expectations family of return forecasts. Second, it is closely related to other familiar conditional variance estimates.¹⁷ Finally, the gain level for the variance in a forecast rule, $g_{j,\sigma}$, is allowed to be different from that used in the mean expectations, g_j . This allows for rules to have a different time series perspective on returns and volatility.

2.6 Market Clearing

The market is cleared by setting the individual share demands equal to the aggregate share supply of unity,

$$1 = \sum_{i=1}^I Z_{t,i}(P_t). \quad (24)$$

Writing the demand for shares as its fraction of current wealth, remembering that $\alpha_{t,i}$ is a function of the current price gives

$$P_t Z_{t,i} = (1 - \lambda) \alpha_{t,i}(P_t) W_{t,i}, \quad (25)$$

$$Z_{t,i}(P_t) = (1 - \lambda) \alpha_{t,i}(P_t) \frac{(P_t + D_t) S_{t-1,i} + B_{t-1,i}}{P_t}. \quad (26)$$

This market is cleared for the current price level P_t . This needs to be done numerically given the complexities of the various demand functions and forecasts, and also the boundary conditions on $\alpha_{t,i}$.¹⁸ It is important to note again, that forecasts are conditional on the price at time t , so the market clearing involves finding a price which clears the market for all agent demands, allowing these demands to be conditioned on their forecasts of R_{t+1} given the current price and dividend.¹⁹

2.7 Gain Levels

An important design question for the simulation is how to set the range of gain levels for the various forecast rules. These will determine the dynamics of forecasts. Given that this is an entire distribution of values it will be impossible to accomplish much in terms of sensitivity

¹⁶See Nelson (1992) for early work on this topic.

¹⁷ See Bollerslev, Engle & Nelson (1995) or Andersen, Bollerslev, Christoffersen & Diebold (2006) for surveys of the large literature on volatility modeling.

¹⁸A binary search is used to find the market clearing price using starting information from P_{t-1} .

¹⁹ The current price determines R_t which is an input into both the adaptive, and short AR forecasts. Also, the price level P_t enters into the P_t/D_t ratio which is required for the fundamental forecasts. All forecasts are updated with this time t information in the market clearing process.

analysis on this. Therefore, a reasonable mechanism will be used to generate these, and this will be used in all the simulations.

Gain levels will be thought of using their half-life equivalents, since the gain numbers themselves do not offer much in the way of economic or forecasting intuition. For this think of the simple exponential forecast mechanism with

$$f_{t+1}^j = (1 - g_j)f_t^j + g_j e_{t+1}. \quad (27)$$

This easily maps to the simple exponential forecast rule,

$$f_t = \sum_{k=1}^{\infty} (1 - g_j)^k e_{t-k}. \quad (28)$$

The half-life of this forecast corresponds to the number of periods, m_h , which drops the weight to $1/2$,

$$\frac{1}{2} = (1 - g_j)^{m_h}, \quad (29)$$

or

$$g_j = 1 - 2^{-1/m_h}. \quad (30)$$

The distribution of m_h then is the key object of choice here. It is chosen so that $\log_2(m_h)$ is distributed uniformly between a given minimum and maximum value. The gain levels are further simplified to use only 5 discrete values. These are given in table 1, and are $[1, 2.5, 7, 18, 50]$ years respectively. In the long memory (low gain) experiments these five values will be distributed between 45 and 50 years.

These distributions are used for all forecasting rules. All forecast rules need a gain both for the expected return forecast, and the variance forecast. These will be chosen independently from each other. This allows for agents to have differing perspectives on the importance of past data for the expected return and variance processes.

2.8 Adaptive rule selection

The design of the models used here allows for both passive and active learning. Passive learning corresponds to the long term evolution of wealth across strategies. Beyond passive learning, the model allows for active learning, or more adaptive rule selection. This mechanism addresses the fact that agents will seek out strategies which best optimize their estimated objective functions. In this sense it is a form of adaptive utility maximization.

Implementing such a learning process opens a large number of design questions. This

paper will stay with a relatively simple implementation. The first question is how to deal with estimating expected utility. Expected utility will be estimated using an exponentially weighted average over the recent past,

$$\hat{U}_{t,j} = \hat{U}_{t-1,j} + g_u^i(U_{t,j} - \hat{U}_{t-1,j}), \quad (31)$$

where $U_{t,j}$ is the realized utility for rule j received at time t . This corresponds to,

$$U_{t,j} = \frac{1}{1-\gamma}(1 + R_{t,j}^p)^{(1-\gamma)} \quad (32)$$

with $R_{t,j}^p$ the portfolio holdings of rule j at time t . Each rule reports this value for the 5 discrete agent gain parameters, g_u^i . Agents choose rules optimally using the objective that corresponds to their specific perspective on the past, g_u^i , which is a fixed characteristic. The gain parameter g_u^i follows the same discrete distribution as that for the expected return and variance forecasts.

The final component to the learning dynamic is how the agents make the decision to change rules. The mechanism is simple, but designed to capture a kind of heterogeneous updating that seems plausible. Each period a certain fraction, L , of agents is chosen at random. Each one randomly chooses a new rule out of the set of all rules. If this rule exceeds the current one in terms of estimated expected utility, then the agent switches forecasting rules.

3 Results

3.1 Benchmarks

This section presents results from two sets of benchmark parameters for the model. The first restricts gain parameters to low values corresponding to long learning horizons. It will often be referred to as the “low gain” experiment. This is an important test to see if the model is capable of converging to a reasonable steady state, and whether the learning dynamics will function as long as all gain levels are set low enough, and agents are using data over long ranges in the forecasts. The half-life ranges for gain parameters in these experiments are set to 45 – 50 years. Also, for this run only the coefficient of relative risk aversion is raised from 3.5 to 8. This is necessary since the model converges to a very low return variance. This would drive all portfolio holdings to their extreme, α_H , in this situation which makes the market difficult to clear. The model is run for 200,000 weeks to make

sure it has reached some form of long run steady state, and all early transients have been eliminated. Figure 1 displays the price/dividend ratio, returns, and trading volume from the last 100 years of simulated data. Prices do not move far from dividends as shown in the top panel. Returns appear to be regular, and they can be shown to be uncorrelated, and close to Gaussian. Trading volume is nearly constant, reflecting the small amount of noise added to each agent's portfolio demand.

This simulated set of prices and returns generates an important, but relatively uninteresting set of results from the market simulation. While these features are not close to those of actual markets they do demonstrate that for a certain set of parameters this market and the agents' learning algorithms are capable of converging to a reasonable approximation to a rational expectations equilibrium. The most important aspect of this set of parameters is the fact that agents use learning algorithms which force them to concentrate on long range averages as opposed to short run trends. Also, though the aggregate data looks like an equilibrium, it is not quite an equilibrium in the traditional sense. There is still trading, and shifts in strategies which do not appear in the aggregate time series.

Figure 2 displays results from a run that will be used extensively to demonstrate the features of heterogeneous gain learning in this model. In this case, the 5 gain levels used vary from 1 year to 50 years as shown in table 1. This will often be referred to as the "all gain" experiment. The three panels of figure 2 should be compared with the previous figure. Now the price/dividend ratio varies erratically through time, which is important to generating reasonable long swings in asset prices. The weekly returns demonstrate two features common to relatively high frequency returns. First, there are a large number of extreme returns, and second, periods of high volatility appear clumped together. These last two facts are interesting, and are important to the dynamics that this model generates. However, they are secondary to the main mission of this paper which is concerned with the top panel, and long swings from fundamentals.

The large swings in the price/dividend ratio are compared to the values from the actual series in figure 3. The persistent, but erratic swings in the actual P/D ratio are clear.²⁰ This is not an easy series to characterize, and it is possible that visual analysis may be the best that can be done. It would appear that there are fluctuations that occur at frequencies which are large relative to the length of the series, and one could also question the stationarity of this series as well. The long gain half-life is set to 50 years in recognition that this length of data (140 years) is probably the best that many investors can do in accessing the long

²⁰ The S&P series used here is the Shiller annual series available from <http://www.econ.yale.edu/shiller/>. This series contains both dividend and earnings information. The rough visual features of this series would be similar if the price/dividend ratio were replaced with the annual price/earnings ratio.

range dynamics for stock returns. The panels below the actual data represent 3 random snapshots from runs of the simulations using the long and short gain levels. They draw 140 year periods from a simulated time series of 4,000 years, so there is little chance for overlap. These visually show similar patterns to the actual data with varying levels of persistence and variability. It is interesting to note that all series are capable of long periods of near constant ratios. Eyeball tests of time series should always be viewed with suspicion, but in this case the special nature of the target series, and the sorts of dynamics we are interested in makes these tests essential.²¹

Table 2 presents summary statistics on annual data comparing the all gain level simulations with the annual Shiller data. Simulation results are taken from the final 1,000 years in the 4,000 year all gain run. Comparisons are made to both the P/D ratios and the P/E ratio since the model could be viewed as generating either of these. The table shows that the simulated data generates reasonable levels and variability in the first two rows. It is a little low relative to the mean P/D ratio, but it should be realized that much of that large mean, 26, is driven by observations of the late 1990's or early 2000's. The model does generate a slightly lower level of variability with a standard deviation of only 4.24 which is smaller than the two ratio comparison series. There is some evidence of this smoother dynamics in figure 3, but it should be stressed that all of these moments from the data are estimated with very low precision.

Table 2 also reports the annual real return and standard deviation for the simulation and the actual data along with the Sharpe ratio. These are comparable, but the simulation generates a slightly higher level of return variability with an annual standard deviation of 0.25 as compared to the historical value of 0.17. This is common for many simulation runs. Two issues should be kept in mind when comparing these values. First, the data moments are again not estimated with great precision, and 0.25 would fall within a reasonable range of 0.17. More importantly, the simulation is concerned with generating the dynamics of a single representative asset. Given this, it is not clear if the variability of its return output should be judged in comparison to the index, or to the much higher level of volatility displayed by individual stocks. Table 2 shows that the simulations are in a reasonable range when compared with long range data.

The most demanding and controversial aspect of long range deviations from fundamentals is that it generates some observed predictability in long range returns. This is explored in table 3 which implements standard long horizon forecasting regressions of 1 year and 5 years on the current logged P/D ratio. This is done both for the all gain simulation, and the low

²¹The skeptic could argue that the original series is simply too nonstationary to deal with. However, even this statement makes a big point about learning and long range financial data. What should investors do in that case?

gain only benchmark which was displayed in figure 1. In all cases the dependent variable is the log return at either the 1 or 5 year horizon. These returns are regressed on both the $\log(P/D)$ ratio and also the $\log(C/D)$ ratio. This latter ratio is an important consumption ratio for this model since consumption can only be funded through dividend flows. When this ratio is less (greater) than 1 agents are increasing (decreasing) their savings through changes to their asset positions. Row one shows that the regression is capable of forecasting future returns with a highly significant coefficient, and a R^2 of nearly 0.19. Relative to actual data this is large, but it should be remembered that this regression is run on a sample length of 1,000 years.²² The second row adds the consumption ratio which is highly significant, but does not add significantly to the R^2 . The third row reports the regression for the 5 year horizon. The coefficient is still highly significant, but the R^2 falls to 0.14.

The lower panel of table 3 reports the results for the small gain (long half-life) only rules. The results for these simulations were displayed as the important initial benchmark in figure 1. In all cases the predictability is near zero. All of the regression coefficients are insignificant with near zero R^2 . This model again displays features which are consistent with an efficient market driven by well functioning simple learning algorithms. They have achieved their mission of driving predictability from the market.

The next two figures deal with shorter horizon phenomenon generated in the all gain simulations which have important implications for the longer horizon returns. Figure 4 compares the distribution of weekly returns for both the CRSP value weighted index from 1926-2009, and the all gain simulation shown in figure 2. These are drawn from the final 50,000 weeks of a 200,000 week simulation. Both distributions show that they are not Gaussian, and are leptokurtic, possessing “fat tail” features. This result is well known for high frequency asset returns in most markets. While this is not a long range feature, the high frequency of extreme moves in these distributions does impact the learning dynamics of the various strategies.

Figure 5 summarizes two key time series features of the generated returns, and compares these with the CRSP benchmark. The top panel reports return autocorrelations for the CRSP index, and the same sample of 50,000 weeks used in the previous figure. Both series show very little evidence for any amount of correlation in returns. The lower panel displays the weekly autocorrelation for absolute returns for the same data. The simulation generates very strong positive autocorrelations as do the actual returns. However, the correlation

²² It would be interesting to see how well this works on sample sizes comparable to the real data series. In these series the issue of the usefulness of predictability is still an open question. See Goyal & Welch (2003) for some of the debate on out of sample predictability, and Fisher & Statman (2006) for some skeptical views on the usefulness of trading rules based on P/D ratios.

pattern is a little high at smaller lags, and drops to zero by the lag at 100 weeks.²³ Though this feature is short range in nature, its long correlation pattern demonstrates that it may be important for long range patterns in financial data. Further evidence in support of this will be presented in the next section.

3.2 Wealth Distributions

The previous section demonstrated that the benchmark simulation model with forecast rules and agents using small and large gain level learning algorithms can generate reasonably calibrated time series. These series are not perfect, but still generate features that are hard to replicate in many traditional theoretical frameworks. Generation of these features is interesting, but is only a small part of model construction. Do these models give any insights into what is driving these results? Can one look into their inner workings to better understand the features from the actual data? This section will explore this issue in terms of the initial question about long swings of prices from fundamentals using the all gain simulation as a test bed for this.

Beyond prices and returns, the model provides a rich set of information about the dynamics of the forecasting rules, the learning algorithms, and wealth distributions. How wealth endogenously falls across forecasting strategies is the key state variable in the model. This determines the dynamics of prices and trading volume as the market moves, but it is important to remember that these wealth distributions are moving targets as they change through time as well. Figure 6 presents the time series of wealth fractions for the 4 forecast families over the entire 200,000 week ($\sim 4,000$ year) period. These will vary across runs, but this picture is generally representative of many other simulations. There are many important features in the figure. First, the wealth fractions do vary over time. There is no tendency to settle down to relatively constant wealth fractions. The fractions appear to take swings that can both be short and long lived. Similar to our initial figures on long swings in P/D ratios this figure shows swings in wealth ratios that can last for almost 100 year periods.

While the time series are not completely smooth, they do appear to converge to a relatively stable ranking. The simple buy and hold strategy, which invests a constant fraction in stock based on the long range forecast for the expected return and variance, controls roughly 50

²³ Modeling the dynamics of volatility at longer horizons is far from settled in finance. For these figures one would see dramatic instability in the autocorrelation pattern depending on whether the Great Depression was added to the series or not. Also, several authors have proposed formal fractionally integrated models such as the FIGARCH, Baillie, Bollerslev & Mikkelsen (1996). Models which suggest multiple (but finite) numbers of horizons have also been suggested in LeBaron (2001*b*) and Corsi (2009). Some authors have also proposed using splines to fit these low frequency components as in Engle & Rangel (2008). Given what we know, one could say that the model and the data both generate a form of long range persistence similar to that observed in the fundamental ratios.

percent of wealth in the market over the entire range. This might be surprising in that this passive strategy cannot contribute to dramatic fluctuations in prices, but it is important to remember that in a heterogeneous world, this type may not be the marginal investor, critical for pricing.²⁴ Also, it is interesting that any model generating reasonable financial time series should also generate evidence which is at least roughly consistent with efficient markets at some horizons, and therefore the objective functions of many investors will find the passive strategies optimal. Basically, the gains to dynamic trading strategies do not dramatically dominate the passive strategies, and in a noisy decision making world some investors will go with each of these.

The next two strategies by wealth in figure 6 are the adaptive, trend following family, and the fundamental family. The battle between these strategies is critical to the overall stability of the market. Through the simulation the adaptive types control nearly 30 percent of wealth as compared with 15 percent for the fundamental types. This larger share of wealth in what are essentially destabilizing strategies is important to the behavior of the market at both long and short horizons, and some examples will be given in later figures.²⁵ Finally, the short AR traders, who use a short range regression to forecast returns, remains a small part of the market at about 5 percent of wealth. This strategy follows a classic efficient market dynamic. Its presence drives short range return correlations to zero while putting itself nearly out of business. It always stays in the fringe, soaking up some short term predictability, and ready to take action anytime large short term correlations appear, but it never takes a large share of wealth.²⁶

The strategy families are only one broad distinction over the forecast rules. Each strategy in a family is also characterized by its forecast gain parameter, and its variance forecast gain parameter. As mentioned previously, these are critical in determining how the agents weigh past data in their forecasts. It has already been shown that limiting these two parameters to small values across all forecast rules leads to convergence of the market into a rough approximation of a rational expectations equilibrium. With lots of gain levels active, more realistic price and return series are generated. Figure 7 shows the distribution of wealth across the 5 discrete gain levels. These are the same half-lives as shown in table 1 and would correspond (left to right in the figure) to [50, 18, 7, 2.5, 1] years. For both the adaptive and fundamental strategies, wealth is not converging to any particular gain level, and, most importantly, there is no strong tendency for wealth to move to the smallest possible gain

²⁴See Adam & Marcet (2010b) for some recent theoretical results on heterogeneous agents and pricing.

²⁵The basic fact that the fundamental strategies may not control enough wealth to mitigate extreme moves in asset prices is related to “limits to arbitrage” ideas as in Shleifer & Vishny (1997).

²⁶Simulation experiments without this strategy show that its role is critical. Removing it generates time series with short range autocorrelations much larger than actual financial series.

(largest half-life) strategy. The situation is different for the short AR strategy which shows wealth moving to the smaller gain strategies. In this case, almost 75 percent of wealth is using forecasting models with gain half-lives of 18 or 50 years which seems comfortably long for estimating weekly regression models. The buy and hold strategies only use long half-life forecasts in their rules, so their distributions are not presented.²⁷

Figure 8 reports the results for the variance forecast gain levels. In this case the figure shows a more distinct pattern with a large amount of wealth concentrated on the large gain (low half-life) forecasts. The rules still leave a moderate amount of wealth at longer horizons, but in each case there is a distinct concentration of wealth in rules which utilize a small amount of data in their volatility forecasts.

In the model a key gain parameter comes in how agents assess the value of the different available strategies. This was described in equation, 31. Figure 9 reports the distribution of wealth across utility gain levels g_u^i . The gain level is considered a fixed component of all agents beliefs about the world. This plot shows whether there is any tendency for any type, particularly those using longer horizons of data, to dominate the market. This does not appear to be the case, and all types survive. There is some weak evidence for some of the lower gain, longer half-life, strategies to do better, but it appears to be very weak.

Wealth distributions show a few broad features across the model. Buy and hold strategies dominate, but adaptive strategies continue to hold a large amount of wealth as well. Fundamental strategies reveal a significantly smaller fraction of the market relative to these other two strategies which is important in market dynamics. In terms of gain levels, in most cases there is no strong selection for low gain learning models which would use large amounts of data. This is important since it is this restriction that allows the learning models to converge as shown in the first benchmark simulation. The next section will explore some of the long range dynamics given this population of agents and rules.

3.3 Mechanisms

Since the model gives full access to the strategies, it is insightful to look inside and see exactly what they are doing. The next few figures produce some snapshots of prices and forecasts which show the general pattern of forecast behavior as it responds to price movements. Figure 10 displays the dynamics of the wealth weighted expected returns across all four forecast families along with the price level. This is done over a 3000 week (or roughly 60 year) snapshot of data. The second panel shows the forecast for the two most important

²⁷These histograms are generated as a single period snapshot at the end of the 200,000 week run. They should be viewed with some caution since sampling variability over time could be important. Specific patterns may not be as strong as they appear, but the basic idea that for two of the forecasts, many gain levels remain active in the market is generally robust.

strategies, adaptive and fundamental. These move as they should, and they are dominated by their large moves near big price drops. The adaptive forecasts observe the down trend in the market, and make low return forecasts while the fundamental forecasts see the price as being low relative to the fundamentals, and respond with a high conditional return forecast. Qualitatively, this is as it should be for these strategies.

The lower panel reports the results for the buy and hold strategy, and the short AR strategy. The buy and hold is basically a constant, corresponding to its long horizon mean. The short AR strategy generates a lot of variability around this value with high volatility, and large swings. Since this strategy is driven purely as an autoregressive model on lagged returns it is sensible that the forecast itself picks up much of the return variability. This intense variability in the strategy may be related to its rather poor performance.

Figure 11 presents the same time series snapshot, but now replaces the expected return forecast with the actual portfolio strategy for each family as represented by the equity fraction. This involves combining the expected return forecast along with the variance forecast. This figure shows some more interesting dynamics than the last one. The top panel shows the price levels, and the middle and lower panels are the different forecast family strategies. First, the adaptive forecasts do behave as expected in following general price trends. They appear to be readjusting portfolio weights very quickly in response to small price trends, and they also back off their positions substantially after large price declines. The behavior of the fundamental strategy is more interesting. During large price increases it actually starts increasing its holding of stock. This is curious since it doesn't make sense for this strategy as the price is moving farther out of line with dividends. We will see that there are two mechanisms at work here. First, during major price run ups volatility falls, so their estimated risk levels are falling, pushing them to a more aggressive portfolio regardless of their conditional return forecasts. Also, these forecasts are dependent on estimates of their P/D regressions. As a price rise continues, evidence in support of reversion to the fundamental diminishes in strength. In other words, they begin to lose faith in their fundamental models, and this would sweep across their populations from small to large half-lives. It is also true that their response to a large price drop seems surprisingly weak. This is also a reminder that their behavior is impacted by the conditional variance. While their expected returns are high after a fall, these large returns are feeding into their variance models, and they believe that risk is high in these states of the world. They are nervous and may not take as aggressive a position as their conditional return models suggest.

The features of the fundamental trading strategies become clearer when they are analyzed over the extreme gain levels as in figure 12. The middle panel in this figure looks at the

fundamental strategies for the highest and lowest gain level in the volatility component only. If the low gain strategies are relatively long term in their assessment of risk, they should respond to a fall in price more aggressively than the short gain volatility strategies. This is very dramatic in the figure. Low gain strategies come into the market immediately on a price fall, and max out their portfolio while the high gain strategies do not behave as aggressively.²⁸ The bottom panel presents the estimates of their least squares coefficient from the return regressions on the value for $\log(P/D)$. For the low gain strategies this is stable and negative.²⁹ The high gain forecasts show a pattern which moves through time in a predictable fashion. It maxes out at 0, which is its maximum allowable value, during significant price run ups. The short range of data that this regression is relying on shows no strong return to fundamentals, and the regression coefficient simply reflects this belief. After a large price drop, the parameter swings negative, again reflecting that the recent data points now strongly support the fact that large negative returns often follow periods when prices are large relative to fundamentals. If these results seem somewhat extreme, it is important to remember that these are the extreme gain levels. There are 3 others which lie in between these. The aggregate forecast depends on how wealth falls across these different strategies as given in figure 7.

The next figures try to explore the mechanism behind the large swings in fundamentals in the model. Specifically, the hypothesis of whether the swings are driven by the common component in agents' forecast models, their prediction of the conditional variance. Figure 13 is a simple plot of the $\log(P/D)$ ratio versus the log of the wealth weighted variance forecast in the market. The negative correlation is very strong, and a regression of $\log(P/D)$ on the logged variance forecast yields an $R^2 = 0.72$.³⁰

Further evidence on this connection is given in the time series results in figure 14. The top panel plots the P/D ratio for a 50 year sample of time, and the middle panel displays the wealth weighted variance forecast. The wealth weighted variance clearly moves closely with the P/D ratio. Large price falls cause a jump in the variance estimates, and these generate a persistent increase in the market average variance forecast.³¹ Now the issue of

²⁸ This hesitancy on the part of risk averse rational investors is in the spirit of noise trader models as in DeLong, Shleifer, Summers & Waldman (1992).

²⁹ These regressions are run on the mean weekly return over 52 weeks as opposed to annual returns themselves. This is why the parameter differs from those in table 3. Multiplying by 52 shows that they are indeed consistent with each other in magnitude.

³⁰ Obviously, this right hand side variable here is endogenous, so this regression should only be viewed as instructive about the comovements.

³¹ A quick back of the envelope test is useful to understand the impact of this volatility change on prices. A rough representative agent approximation to the pricing relation would give, $P = (\frac{x}{\gamma\sigma^2} + \frac{1}{2\gamma})W$. Ignoring the 2nd term in the sum as being relatively small yields a direct proportionate connection between the price level and the variance for a given wealth level. An increase in σ^2 by a factor of 3 as shown in the figure could lead to a reduction in prices by the same factor which is a similar order of magnitude as shown in the figure. This reduction would be larger if one correctly considered the decrease in wealth as

which variance forecast is being used becomes important. The lower panel in the figure displays the variance forecasts for the highest and lowest gain levels across all forecast rules. The low gain (or long half-life) forecast generates a very stable estimate of the conditional variance. It is remarkably stable even after the occurrence of a large price drop. On the other hand, the high gain forecast moves quickly with the new information coming from the price fall. The middle panel lies somewhere between these two. It also reveals some evidence for a more persistent impact of volatility, as more of the lower gain strategies pick up the increased volatility levels that are usually persistent. Evidence in this figure is consistent with the portfolio strategy results given in figure 11. The low gain strategies are not changing their risk assessments even after large price declines, and they are therefore very aggressive in their portfolio weights, which other strategies are more cautious about. Higher gain strategies perceived the increase in conditional variance, and this moderates their response to their high conditional return forecasts. In other words, they know it is a good time to invest in equity, but they are scared.

Why don't agents chose what appears to be a more desirable variance forecast? This is an interesting question which is addressed in figure 15. The top two panels in the figure report the mean squared error (MSE) and mean absolute error (MAE) for the different gain level variance forecasts evaluated over the last 50,000 weeks of the simulation. They are normalized by the value from the forecast using the unconditional mean over the sample. A value of 1 would correspond to getting the same prediction error as the unconditional mean. The lowest gain level shows that it is very close to using the unconditional mean squared return which should be expected. What is interesting is that the higher gain forecasts all generate modest improvements in terms of forecast errors. For the highest (lowest half-life) forecast there is an improvement of more than 10 percent for the the MSE and MAE forecast measures. This is important since these are true out of sample forecasts at fixed parameter values. They could have gone well above one in value, and did not have to show an improvement. Also, this result is consistent with the results on wealth distributions over forecasting rules as shown in figure 8. It is not a direct causal path, but it is an indication of why these higher gain forecasting rules seem to be doing well. They are generating an endogenous pattern in volatility which yields some amount of predictability for higher gain learning forecasts over the more stable low gain forecasts, causing these same forecasts to thrive in the population.

These forecast measures are not really relevant to how agents are actually selecting rules which is based on expected utility. The histogram in the lower left corner of figure 15 shows

well.

a snap shot of an expected utility estimate. These rate the overall value of all forecasting rules to agents in certainty equivalent terms across the different volatility gain levels. The figure shows relatively equal values, with some indication for higher utilities for the high gain variance strategy. The dynamic here would be that these strategies more accurately warn agents of the higher volatility periods, allowing them to back off on equity positions during these periods. If this dynamic strategy is effective, then it will be reflected in a higher risk adjusted return for this strategy, which appears to be the case.³²

These results make a strong case for the importance of high gain, short half-life learning rules. The fact that at least some fraction of wealth stays with these rules over time would appear essential to building reasonable market dynamics. However, what about the low gain, long half-life rules? Do they play any role, or are they superfluous in generating reasonable dynamics? Many of the distributions still show a large fraction of wealth using these rules, so the question remains about their impact. To explore this, a reverse of the first benchmark, low gain, experiment is performed. Agents with only high gain, or short half-life rules are used. The set of rules is again concentrated on 5 different gain levels, but instead of being distributed between 1 and 50 years, they are reduced to a range of 1-5 years only. Figure 16 displays a 100 year period of a run with these gain parameters. The market still displays significant instability. However, the dynamics do not appear reasonable for lining up with real data. There are quick bursts in the price level which suddenly take off, and crash almost as quickly. After the sharp drop in prices, prices return relatively quickly to a central P/D level. Returns are punctuated by large tail events, but prolonged periods of high volatility are not evident. The competing dynamics of learning agents of all gain levels would appear essential to spread out some of the market instability, making it less dramatic, and more persistent in all dimensions. This is a critical requirement if one is interested in modeling deviations from fundamentals which are not just large, but are persistent.

4 Conclusions and Implications

This model demonstrates that heterogeneous gain learning may generate reasonable long range dynamics for stock prices. In particular, the model generates persistent price/dividend time series, as well a replicating many other features of asset prices. More importantly, it does so in a framework that uses a simple set of strategies representing behaviors that are probably generic to most financial settings. Finally, when the details of agent behavior

³²The results in this panel should be viewed with some care since they are a one period snapshot. They are probably subject to extensive sampling variation which is not reported in the figure. A cautious interpretation of the figure suggests that utility based models would view all rules equally. However, it does appear that the slightly higher value of the short gain rules is being transmitted into real adjustments of agent wealth toward these strategies.

are analyzed, they generate patterns consistent with individual objective functions. In other words, they appear individually rational, relative to the environment that their own strategies self generate.

The model is still highly stylized, and it would be premature to call this a definitive model for financial markets. It therefore can only be viewed as a computational thought experiment.³³ For example, no attempt is made to calibrate the consumption series, or to model real interest rates. These remain open challenges for the future. The model does quantitatively capture swings in the price/dividend ratio. It does this through the mechanism of the common volatility forecasts, and their endogenous structure across gain levels in their variance forecasts. Obviously, this is the component that one should question in terms of the results. As usual in heterogeneous agent models, questions should be asked about whether these are the types of forecasts used by market participants. Would other variance forecasts change the results? It seems plausible that a wide range of forecasting models will give relatively similar results. One critical issue to consider in the future, is whether the models should entertain some convergence to a long run mean level. The exponential filter is consistent with a world with an unobserved state variable following a random walk as in Muth (1960). This persistence in forecasts obviously contributes to the generated large swings in P/D ratios. It is not clear that adding some kind of long run mean reversion would change the situation, since the measured volatility in the simulation generates a lot of persistence, but this would be an interesting experiment.

Another learning issue that needs to be further explored is how the rules react to large shocks. The market endogenously generates large returns from distributions which are not normal. These have a dramatic impact as inputs into the learning models, and occasionally lead to large swings in parameters which can last for some time. Would real investors blindly feed these extreme values into their own learning algorithms, or would some more complicated filtering algorithm be more appropriate? If the latter is the case, it is still not clear which stylized modeling framework would be appropriate in this context. However, this issue is another one for future consideration, both for learning models in finance and in macro economics.³⁴

This paper has made the case for greater consideration of heterogeneity in gain parameters in learning models. When these models are carefully constructed, they are capable of producing price dynamics replicating most empirical financial market features. High gain learners are a destabilizing force in the market as some agents put heavy weight on the recent

³³The book by Miller & Page (2007) advocates stylized computational experiments which still may not completely capture reality, but influence our intuition about the interactions and dynamics of heterogeneous agent worlds.

³⁴This is area is well underway as in Hansen & Sargent (2008).

past in assessing risk and return. Low gain learners are also essential because their longer term perspective is necessary to generate relatively persistent swings from fundamental values. Finally, when one examines the internals of these models, the micro level behavior appears sensible and consistent, and capable of generating the macro features. The agents self generate a time series world in which their own behavior is reinforced, and a robust range of beliefs survive which continue to perpetuate these patterns. The generated long swings in prices are large, persistent, and unpredictable. All of these features make it difficult for learning agents to lock down on the strategies that would be necessary to eliminate them.

References

- Adam, K. & Marcet, A. (2010*a*), Booms and busts in asset prices, Technical Report 2010-E-2, Bank of Japan, Tokyo, Japan.
- Adam, K. & Marcet, A. (2010*b*), Internal rationality, imperfect market knowledge and asset prices, Technical report, Mannheim University.
- Andersen, T. G., Bollerslev, T., Christoffersen, P. F. & Diebold, F. X. (2006), Volatility and correlation forecasting, *in* G. Elliott, C. W. J. Granger & A. Timmermann, eds, ‘Handbook of Economic Forecasting’, North Holland.
- Baillie, R. T., Bollerslev, T. & Mikkelsen, H.-O. (1996), ‘Fractionally integrated generalized autoregressive conditional heteroskedasticity’, *Journal of Econometrics* **74**, 3–30.
- Blume, L. & Easley, D. (1990), ‘Evolution and market behavior’, *Journal of Economic Theory* **58**, 9–40.
- Bollerslev, T., Engle, R. F. & Nelson, D. B. (1995), ARCH models, *in* ‘Handbook of Econometrics’, Vol. 4, North-Holland, New York, NY.
- Branch, W. A. & Evans, G. W. (2008), Learning about risk and return: A simple model of bubbles and crashes, Technical report, University of California - Irvine, Irvine, CA.
- Brock, W. A. & Hommes, C. H. (1998), ‘Heterogeneous beliefs and routes to chaos in a simple asset pricing model’, *Journal of Economic Dynamics and Control* **22**(8-9), 1235–1274.
- Cagan, P. (1956), The monetary dynamics of hyperinflation, *in* M. Friedman, ed., ‘Studies in the quantity theory of money’, University of Chicago Press, Chicago, IL, pp. 25–120.
- Campbell, J. & Shiller, R. (1988), ‘Stock prices, earnings, and expected dividends’, *Journal of Finance* **63**, 661–676.
- Campbell, J. Y., Lo, A. W. & MacKinlay, A. C. (1996), *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.
- Campbell, J. Y. & Viceira, L. M. (1999), ‘Consumption and portfolio decisions when expected returns are time varying’, *Quarterly Journal of Economics* **114**, 433–495.
- Campbell, J. Y. & Viceira, L. M. (2002), *Strategic Asset Allocation*, Oxford University Press, Oxford, UK.

- Chiarella, C., Dieci, R. & He, X.-Z. (2009), Heterogeneity, market mechanisms, and asset price dynamics, *in* T. Hens & K. R. Schenk-Hoppe, eds, ‘Handbook of Financial Markets: Dynamics and Evolution’, Elsevier, USA, pp. 277–344.
- Corsi, F. (2009), ‘A simple approximate long-memory model of realized volatility’, *Journal of Financial Econometrics* **7**(2), 174–196.
- Dacorogna, M. M., Gencay, R., Muller, U. A., Olsen, R. B. & Pictet, O. V. (2001), *An Introduction to High-Frequency Finance*, Academic Press, San Diego, CA.
- Day, R. H. & Huang, W. H. (1990), ‘Bulls, bears, and market sheep’, *Journal of Economic Behavior and Organization* **14**, 299–330.
- DeLong, J. B., Shleifer, A., Summers, L. & Waldman, R. (1992), ‘Noise trader risk in financial markets’, *Journal of Political Economy* **98**, 703–738.
- Engle, R. F. & Rangel, J. G. (2008), ‘The spline-garch model for low-frequency volatility and its global macroeconomic causes’, *Review of Financial Studies* **21**(3), 1187–1222.
- Evans, G. W. & Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.
- Evans, G. W. & Honkapohja, S. (2008), Leaning in macroeconomics, *in* S. N. Durlauf & L. E. Blume, eds, ‘The New Palgrave Dictionary of Economics’, Palgrave Macmillan.
- Evstigneev, I. V., Hens, T. & Schenk-Hoppe, K. R. (2009), Evolutionary finance, *in* T. Hens & K. R. Schenk-Hoppe, eds, ‘Handbook of Financial Markets: Dynamics and Evolution’, Handbooks in Finance, North-Holland, Amsterdam, the Netherlands, pp. 509–564.
- Fisher, K. L. & Statman, M. (2006), ‘Market timing in regressions and reality’, *Journal of Financial Research* **29**(3), 293–304.
- Frankel, J. A. & Froot, K. A. (1988), Explaining the demand for dollars: International rates of return and the expectations of chartists and fundamentalists, *in* R. Chambers & P. Paarlberg, eds, ‘Agriculture, Macroeconomics, and the Exchange Rate’, Westview Press, Boulder, CO.
- Friedman, M. (1956), *A Theory of the Consumption Function*, Princeton University Press, Princeton, NJ.
- Frydman, R. & Goldberg, M. D. (2007), *Imperfect Knowledge Economics*, Princeton University Press, Princeton, New Jersey.

- Frydman, R. & Phelps, E. S., eds (1983), *Individual Forecasting and Aggregate Outcomes: Rational expectations examined*, Cambridge University Press.
- Fuster, A., Laisson, D. & Mendel, B. (2010), ‘Natural expectations and macroeconomic fluctuations’, *Journal of Economic Perspectives* **24**(4), 67–84.
- Gaunersdorfer, A. (2000), ‘Endogenous fluctuations in a simple asset pricing model with heterogeneous agents’, *Journal of Economic Dynamics and Control* **24**, 799–831.
- Georges, C. (2008), ‘Staggered updating in an artificial market’, *Journal of Economic Dynamics and Control* **32**, 2809–2825.
- Giovannini, A. & Weil, P. (1989), Risk aversion and intertemporal substitution in the capital asset pricing model, Technical Report 2824, National Bureau of Economic Research, Cambridge, MA.
- Goyal, A. & Welch, I. (2003), ‘Predicting the equity premium with dividend ratios’, *Management Science* **49**(5), 639–654.
- Hansen, L. P. & Sargent, T. J. (2008), *Robustness*, Princeton University Press.
- Hommes, C. H. (2010), The heterogeneous expectations hypothesis: Some evidence from the lab, Technical report, CeNDEF, University of Amsterdam.
- Hommes, C. H. & Wagener, F. (2009), Complex evolutionary systems in behavioral finance, *in* T. Hens & K. R. Schenk-Hoppe, eds, ‘Handbook of Financial Markets: Dynamics and Evolution’, North-Holland.
- Hong, H. & Stein, J. C. (1999), ‘A unified theory of underreaction, momentum trading, and overreaction in asset markets’, *Journal of Finance* **54**, 2143–2184.
- Honkapohja, S. & Mitra, K. (2006), ‘Learning stability in economies with heterogeneous agents’, *Review of Economic Dynamics* **9**, 284–309.
- Kelley, J. L. (1956), ‘A new interpretation of information rate’, *Bell Systems Technical Journal* **35**, 917–926.
- Kirman, A. P. (1991), Epidemics of opinion and speculative bubbles in financial markets, *in* M. Taylor, ed., ‘Money and Financial Markets’, Macmillan, London.
- LeBaron, B. (2001a), ‘Evolution and time horizons in an agent based stock market’, *Macroeconomic Dynamics* **5**(2), 225–254.

- LeBaron, B. (2001*b*), ‘Stochastic volatility as a simple generator of apparent financial power laws and long memory’, *Quantitative Finance* **1**, 621–631.
- LeBaron, B. (2006), Agent-based financial markets: Matching stylized facts with style, *in* D. Colander, ed., ‘Post Walrasian Macroeconomics’, Cambridge University Press, pp. 221–235.
- LeBaron, B., Arthur, W. B. & Palmer, R. (1999), ‘Time series properties of an artificial stock market’, *Journal of Economic Dynamics and Control* **23**, 1487–1516.
- Lettau, M. & Ludvigson, S. (2010), Measuring and modeling variation in the risk-return trade-off, *in* Y. Ait-Shalia & L. P. Hansen, eds, ‘Handbook of Financial Econometrics’, Vol. 1, Elsevier, pp. 617–690.
- Lucas, R. E. J. (1978), ‘Asset prices in an exchange economy’, *Econometrica* **46**, 1429–45.
- Lux, T. (1998), ‘The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions’, *Journal of Economic Behavior and Organization* **33**, 143–165.
- Miller, J. H. & Page, S. E. (2007), *Complex Adaptive Systems: An Introduction to computational models of social life*, Princeton University Press.
- Mitra, K. (2005), ‘Is more data better?’, *Journal of Economic Behavior and Organization* **56**, 263–272.
- Muth, J. (1960), ‘Optimal properties of exponentially weighted forecasts’, *Journal of the American Statistical Association* **55**, 299–306.
- Nelson, D. B. (1992), ‘Filtering and forecasting with misspecified ARCH modelsI: Getting the right variance with the wrong model’, *Journal of Econometrics* **52**, 61–90.
- Pastor, L. & Stambaugh, R. F. (2009), ‘Predictive systems: Living with imperfect predictors’, *Journal of Finance* **64**, 1583–1628.
- Pastor, L. & Veronesi, P. (2009), ‘Learning in financial markets’, *Annual Review of Financial Economics* **1**, 361–381.
- Phelps, E. S. (1967), ‘Phillips curves, expectations of inflation and optimal unemployment over time’, *Economica* **34**(135), 254–281.
- Phelps, E. S., ed. (1970), *Microeconomic Foundations of Employment and Inflation Theory*, Norton, New York, NY.

- Sargent, T. (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton, NJ.
- Shiller, R. (1981), ‘Do stock prices move too much to be justified by subsequent changes in dividends?’, *American Economic Review* **71**, 421–36.
- Shleifer, A. & Vishny, R. W. (1997), ‘The limits of arbitrage’, *Journal of Finance* **52**(1), 35–55.
- Tesfatsion, L. & Judd, K. L., eds (2006), *Handbook of Computational Economics: Agent-based computational economics*, North-Holland, Amsterdam.
- Thurner, S., Dockner, E. J. & Gaunersdorfer, A. (2002), Asset price dynamics in a model of investors operating on different time horizons, Technical report, University of Vienna.
- Timmerman, A. (1993), ‘How learning in financial markets generates excess volatility and predictability in stock prices’, *Quarterly Journal of Economics* **108**, 1135–1145.

Table 1: *Parameter Definitions*

Parameter	Value
d_g	0.0126
σ_d	0.12
r_f	0
γ	3.5
λ	0.0007
I	16000
J	4000
g_j	[1, 2.5, 7, 18, 50] years
g_L	50 years
g_u	[1, 2.5, 7, 18, 50] years
L	5 percent/year
$[\alpha_L, \alpha_H]$	[0.05, 0.95]
σ_ϵ	0.02
M_{PD}	52 weeks
M_{AR}	3
h_j	[0.025, 0.15]

The annual standard deviation of dividend growth is set to the level from real dividends in Shiller's annual long range data set. The growth rate of log dividends, 0.0126 corresponds to an expected percentage change of $D_g = 0.02 = d_g + (1/2)\sigma_d^2$ in annual dividends. This corresponds to the long range value of 0.021 in the Shiller data.

Table 2: *Annual Statistics*

	Baseline	Shiller Earnings	Shiller Dividends
Mean(P/D)	15.29	15.37	26.79
Std (P/D)	4.24	5.97	13.90
Autocorrelation(1)	0.64	0.68	0.91
Mean(Log(Return))	6.82		6.22
Std(Return)	0.25		0.17
Annual Sharpe	0.35		0.30

Baseline model uses a sample of 100,000 weeks or about 1900 years. The Shiller series are annual from 1871-2009. P/D refers to the price dividend ratio for the model and the Shiller dividends column. The earnings column uses the annual P/E ratio instead. All returns are real including dividends. The Sharpe ratio estimated from the Shiller annual data uses the 1 year interest rates from that series.

Table 3: *Long Range Return Regressions*

Simulation	$\log(P/D)$	$\log(C/D)$	R^2
All Gain (1 year)	-0.41 (0.03)		0.19
All Gain (1 year)	-0.47 (0.03)	0.15 (0.05)	0.20
All Gain (5 year)	-0.52 (0.09)		0.14
Small Gain (1 year)	-0.01 (0.05)		0.00
Small Gain (1 year)	-0.02 (0.05)	0.06 (0.04)	0.00
Small Gain (5 year)	-0.39 (0.25)		0.01

Dependent variables are 1 and 5 year log simulation returns. Simulation runs use 50,000 weeks of data, or approximately 1,000 years of non overlapping data for regressions. All gain corresponds to runs with both small and large forecast gain levels. Small gain corresponds to runs with only small gains (long memory) forecasts.




Figure 1: **Low Gain Only: Long memory learning**




Figure 2: All Gain : Short and Long Memory




Figure 3: **Price/dividend Ratios: S&P and Three Simulations:** Panel 1 is estimated from the Shiller annual data set and represents the real price dividend ratio for the S&P from 1872-2009. Panels 2-4 are random 140 year snapshots taken from a simulation run of 4,000 years using all gain level learning. The Shiller series goes to a high of near 90 around 2000.




Figure 4: **CRSP VW Index and All Gain Simulation Weekly Return Distributions**




Figure 5: **Return and Absolute Return Autocorrelations**




Figure 6: **Wealth Time Series by Strategy Family**




Figure 7: **Wealth Across Forecast Gain Levels:** The gain levels correspond to half-lives of [50, 18, 7, 2.5, 1] years moving left to right.




Figure 8: **Wealth Across Variance Gain Levels:** The gain levels correspond to half-lives of [50, 18, 7, 2.5, 1] years moving left to right.




Figure 9: **Agent Wealth Across Utility Gain Levels:** The gain levels correspond to half-lives of [50, 18, 7, 2.5, 1] years moving left to right.




Figure 10: Forecast Patterns by Family




Figure 11: **Portfolio Strategies**




Figure 12: **Fundamental Strategies by Gain:** The middle panel shows the equity fractions for the high and low variance gain levels in the fundamental rule family. The lower panel shows the estimated $\log(P/D)$ coefficient for the high and low gain parameters in the recursive least squares regression.




Figure 13: **Volatility Forecast and P/D Ratio:** $R^2 = 0.71$




Figure 14: **Variance Forecast and P/D Ratio Time Series:** “Variance by wealth” is the wealth weighted variance forecast. The lower panel shows the variance forecast averaged (equal weight) across all small and large gain strategies.




Figure 15: **Forecasting and Variance Gain:** MSE and MAE are the mean squared and mean absolute errors for the different variance forecasts estimated across gain levels. They are normalized against the value from guessing the unconditional mean squared return in each case. In other words, 1 indicates no forecast improvement. The utility levels are annualized certainty equivalent returns for the portfolios.




Figure 16: **High Gain Only: Short memory learning:** Simulation results for gain parameters in the range [1-5] years only.